Can Machine Learning help with problems in Quantum Physics?

Optimizing results of quantum experiments with ML

Moritz August, august@in.tum.de
Doctoral candidate at the chair of Scientific Computing, Department of Informatics, TUM
Member of the doctoral programme *Exploring Quantum Matter*

Machine Learning and Quantum Physics might be a nice fit

Quantum Computing → Machine Learning

- Quantum computing holds promise for speeding up ML
- Quantum annealing might help to improve convergence
- Some classical ML algorithms have quantum formulations
- Exponential speed-up in the optimal case
- But many technical issues in the way
- Quantum Annealing not sure to improve convergence speed
- Still, lots to learn by trying

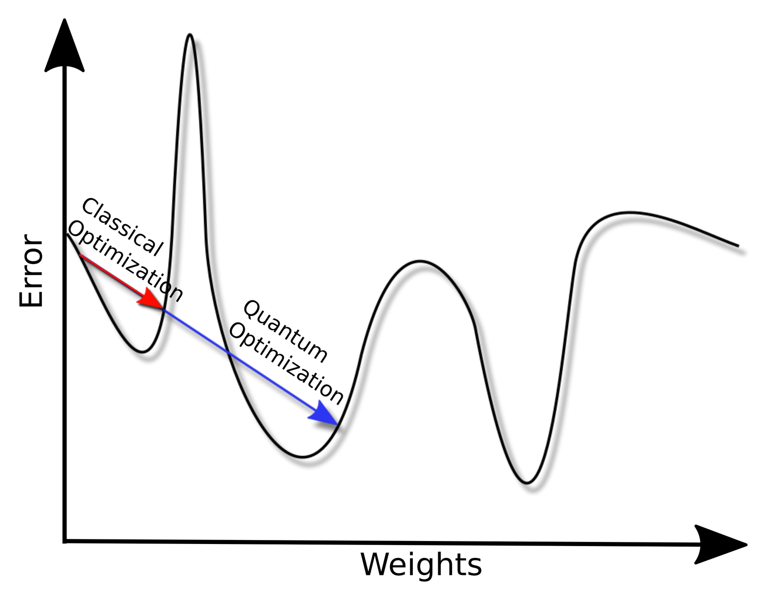

Figure 1: The intuition behind quantum annealing vs. classical optimization in the non-convex case.

Machine Learning → Quantum Computing

- Large-scale quantum computers need yet to be built
- Important problem I: how to preserve qubits?
- Important problem II: how to create quantum gates?
- Both can be perceived as quantum control problems
- ML could help solving such control problems
- More abstractly: Generate optimal/good parameter choices
- Algorithms however need to be compatible with QM

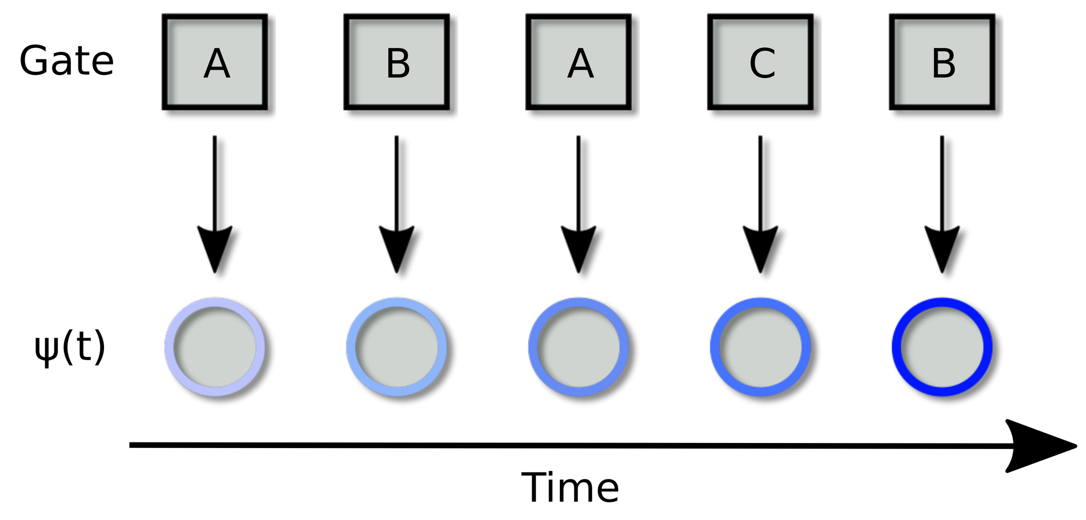

Figure 2: A quantum control problem with a discrete set of possible operations (gates).

Approximating physical properties

- Computing physical properties exactly is hard in QM
- Hilbert space grows exponentially in number of particles
- Sophisticated approximation techniques exist (Tensor Networks)
- Recently, RBMs and CNNs have been used successfully
- Might be alternative way to break the curse of dimensionality
- Interpretability of the models poses a challenge
- Complex weights also necessitate some changes to models

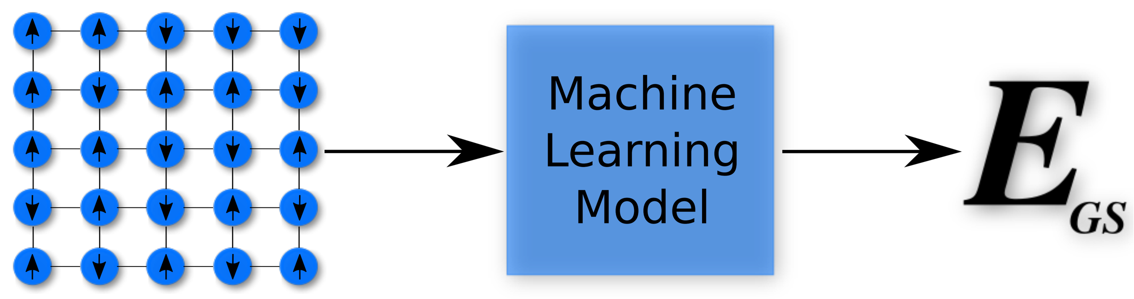

Figure 3: Using machine learning to approximate, e.g., ground state energies of physical systems.

I focus on applying deep learning methods to solve problems in quantum information
An important one is quantum memory

Quantum Memory

- Quantum computers like classical ones must store information
- Qubits are roughly a PMF over the two binary values
- Measurement/Interaction samples from PMF → information lost
- Influence of environment can be seen as noise growing in time
- We need a time evolution $U(t) = e^{iHt)}$ such that for qubit $q(t)$

$$U(T)\cdots U(2)U(1)q(0) \approx q(0), \text{ i.e., } U(T)\cdots U(2)U(1) \approx I$$

in the presence of environment noise

- How to find good/optimal sequences $U(T)\cdots U(2)U(1)$?

Dynamical Decoupling

- DD is a technique to combat the influence of the noise
- Typically, the $U(t)$ are based on a finite set of operations $\mathcal{P}$
- $\mathcal{P}$ commonly is the set of Pauli matrices

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & i \\ -i & 0 \end{bmatrix} \text{ and } Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

and the identity I

- Several analytically derived classes of sequences $\mathcal{P}^{\times T}$ exist
- Optimal under certain noise and strong physical assumptions

Black-Box View

- Supppose we are given a real quantum memory experiment
- Can we find better sequences than analytically known?
- We consider a more realistic setting of assumptions
- Input sequences $s \in \mathcal{P}^{\times T}$ and receive distance measure $\delta(s)$
- The problem to ultimately solve is

$$\min_s \delta(s), \delta(s) \text{ incorporating environment noise}$$

- Known DD sequences exhibit strong local structure
- LSTMs should be good at learning that structure

The evolutionary approach

Algorithm

- Idea: Alternatingly approximate data generating distribution and generate better data
- Assumption: Good sequences have some common structure, no uniform distribution
- For a data set of fixed size D and a model with weights θ, alternatingly solve

$$\min_D \langle \delta(s) \rangle_D \text{ and } \max_\theta \mathcal{L}(\theta|D),$$

- We assume $s_t \sim Cat(s_{t-1}, \ldots, s_1, \theta)$, hence the NLL of a sequence s is given by

$$\text{NLL}(s, \theta) = -\sum_t \sum_i [s_t = i] \log p_i(s_{t-1}, \ldots, s_1, \theta)$$

- $p(s_{t-1}, \ldots, s_1, \theta)$ is represented by an LSTM and the above NLL is minimized to find θ
- In fact multiple LSTMs are trained in every generation to increase robustness

Results

- The algorithm does in fact find sequences better than the ones analytically derived
- LSTMs seem to be able to learn the structure of good sequences

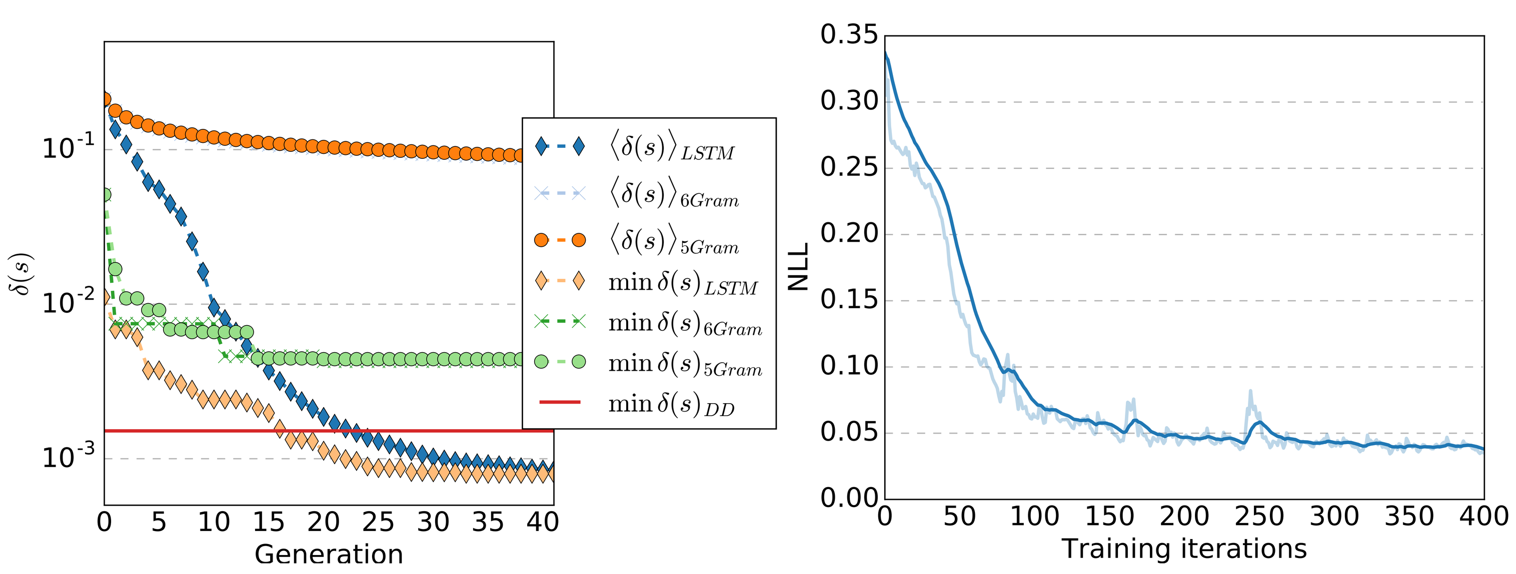

Figure 4: Left: Comparison of the convergence of average and minimal $\delta(s)$ over the number of generations between LSTM and N-Gram models. Right: Convergence of LSTM training in the NLL.

The policy gradient approach

Algorithm

- Quantum experiments and reinforcement learning in principle don't fit so well
- No measurement during the experiment means no rewards, only one reward at the end
- However we can use the REINFORCE method to obtain

$$\nabla_\theta \mathbf{E}[\delta(s)] = \mathbf{E}[\nabla_\theta \log p(s|\theta)(\delta(s) - b)] = \mathbf{E}[\sum_t \sum_i [s_t = i] \nabla_\theta \log p_i(s_{t-1}, \ldots, s_1, \theta)(\delta(s) - b)]$$

- We can approximate the expectation value by generating many sequences (is physically feasible)
- To speed up convergence, I actually use a somewhat modified version of the original distance

$$\delta(s) = \sqrt{1 - \frac{1}{d_S d_B} \|\text{Tr}_S(U(s))\|_{\text{Tr}}}$$

Results

- The reinforcement learning algorithm is able to reproduce results of the evolutionary method
- Much more efficient then the evolution-style algorithm as only one model needs to be trained

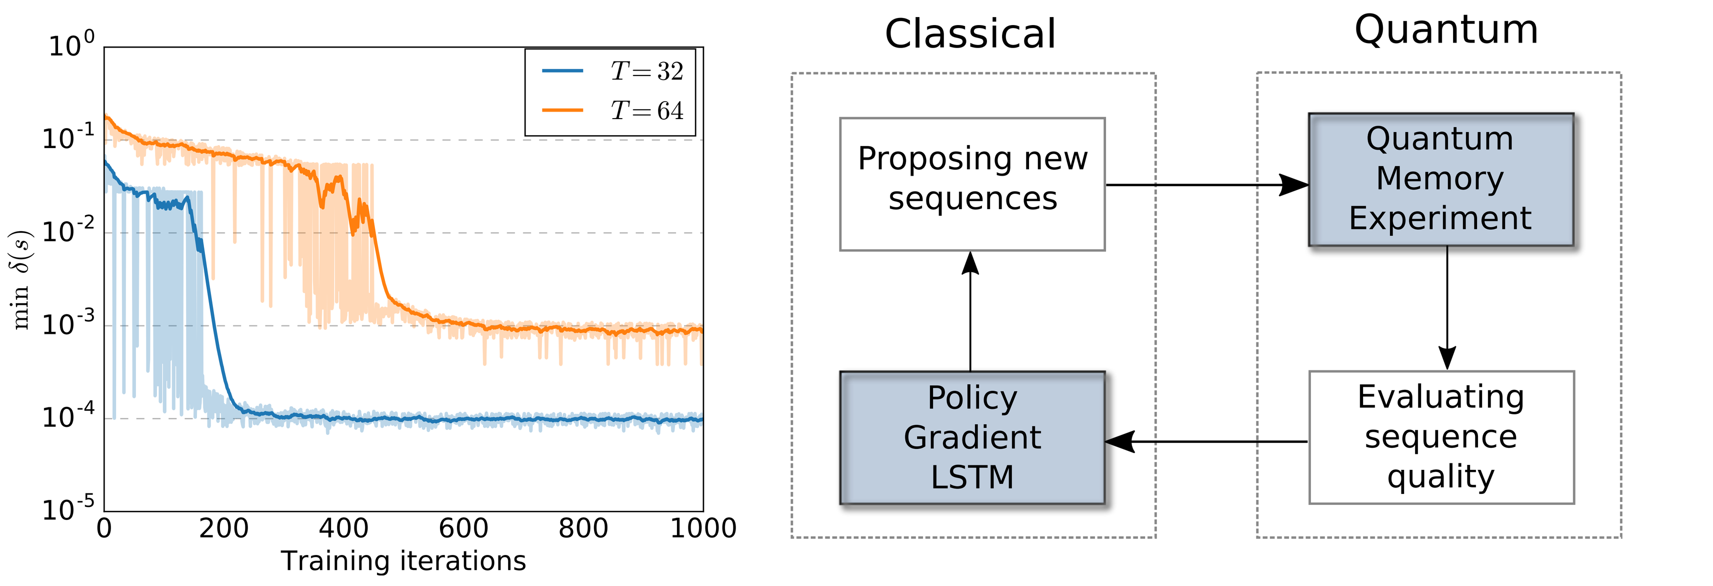

Figure 5: Left: Convergence of the best generated sequences over the training iterations between two sequence lengths. Right: The interaction between the RL agent and the quantum memory.

[1] Moritz August. Quantum-compatible reinforcement learning applied to quantum memory optimization. *in preparation*, 2017.
[2] Moritz August and Xiaotong Ni. Using recurrent neural networks to optimize dynamical decoupling for quantum memory. *Phys. Rev. A*, 95:012335, Jan 2017.
[3] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *arXiv preprint arXiv:1611.09347*, 2016.
[4] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
[5] Juan Carrasquilla and Roger G Melko. Machine learning phases of matter. *Nature Physics*, 2017.
[6] A.B. Finnila, M.A. Gomez, C. Sebenik, C. Stenson, and J.D. Doll. Quantum annealing: a new method for minimizing multidimensional functions. *Chemical physics letters*, 219(5-6):343–348, 1994.
[7] A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel. Active learning machine learns to create new quantum experiments. *arXiv preprint arXiv:1706.00868*, 2017.